Exhibit A

UNITED STATES DISTRICT COURT

DISTRICT OF NEW JERSEY

IN RE: VALSARTAN N-NITROSODIMETHYLAMINE (NDMA) CONTAMINATION PRODUCTS LIABILITY LITIGATION	HON. ROBERT B. KLUGER Civil No. 19-2875 (RBK/JS)

DECLARATION OF JONATHAN JAFFE

I, Jonathan Jaffe, declare, under penalty of perjury, pursuant to 28 U.S.C. § 1746, as follows:

I, Jonathan Jaffe, am the founder and owner of Its-Your-Internet, a technology and ESI discovery consulting firm, established in 2008.

As the Court is aware, I have assisted Plaintiffs in this litigation on ESI discovery issues, including more recently the responses to the proposal of TEVA to utilize what they have termed “CMML” or what is referred to in the industry as Continuous Active Learning “CAL” or “TAR 2.0” in their review process.

I was part of the July 8th telephonic meet and confer with TEVA and their ESI experts. I attended by telephone the hearing held by the Court on July 15th.

I have also been privy to all the communications via letter and e-mail between the parties as well as submissions to the Court regarding this matter.

For the past fifteen years, I have provided ESI discovery and litigation consulting for attorneys in complex litigation throughout the lifecycle of each litigation (from inception to the ultimate resolution). In multiple multi-district litigations inclusive of this one, I have assisted with initial efforts by the parties to reach agreement on search methodologies, as well as to discover the structure of Defendants' systems for maintaining ESI and physical materials. In that capacity I worked with Plaintiffs' counsel in the In Re Benicar (Olmesartan) Products Liability Litigation that was previously litigated before this Court.

I graduated magna cum laude from Columbia University with a BA in Economics-Mathematics in 1999.

For the Court's reference, a copy of my *curriculum vitae* is attached hereto as **Exhibit A**.

While I have been retained by the Plaintiffs, the opinions I express below are independent of any party.

ALLOWING TEVA TO USE TAR/CMML/CAL TO PRIORITIZE AND FILTER DOCUMENTS FOR PRODUCTION IS HIGHLY PREJUDICIAL TO PLAINTIFFS

At the hearing on July 15th, and in submissions to the Court, TEVA has made multiple representations as to the nature of CAL, specifically as to what it can do and what it cannot. As set forth in more detail below, there are fundamental flaws in TEVA's approach, including for example fundamentally incorrect assumptions underlying the arguments as to the value of this system for this particular review for production, the failure to faithfully work to satisfy the Plaintiffs' prioritization requests, the proposed application of the CAL system to the truncated document set already narrowed by the search terms, rather than to the entire document set as is intended, TEVA's refusal to engage in a transparent discussion and collaboration designed to establish the reasonable tagging decisions to apply at the outset, and TEVA's refusal to involve the Plaintiffs throughout the process to ensure faithful application of the tags and continuous quality control and adjustments to continually guide the learning process of the system.

IMPROPER APPLICATION OF CAL AFTER SEARCH TERMS

In the search term negotiation process in which I was heavily involved, Plaintiffs negotiated with the understanding that TAR and CAL would not be used in the review and production process, as no Defendant indicated that it would be using TAR or CAL for review and production.

If a CAL or TAR process is used, it should be applied across the entire set of collected documents. That is how the systems, including the system at issue here, are best utilized. See: "White Papers" attached to the Plaintiffs' July 14, 2020 submission.

Plaintiffs agreed to search term modifications that were only reasonable if TAR and CAL were not going to be used. If this intent had been disclosed, my strong advice would have been not to agree to the modifications and that the present objections would have been presented at that time.

TEVA'S ARGUMENT IS BASED ON A FALSE PREMISE

Notably, in their late evening submission on July 14th to the Court, the TEVA Defendants wrote the following in footnote #1:

> CMML is a technology platform that utilizes a standardized, defensible process to analyze textual content to allow for rapid relevant document discovery. The predictive scores from the CMML model are used to prioritize document review, ensuring that the most relevant documents are reviewed (and produced) first.

Defense counsel repeated this assertion that the prioritization afforded by CMML would result in Plaintiffs receiving the **most relevant documents** in a prioritized manner (though never specifically explaining how each request for prioritization would be fulfilled), during the July 15th, 2020 conference with the Court.

These statements are not correct. Counsel and TEVA's statements are premised on an inaccurate assumption as to a link between the relevancy scores and the *degree* of relevance.

The predictive scores (usually expressed in a range between zero and one) outputted by a CAL system are a representation of how well matched the machine believes a particular document is to the set of documents previously tagged.

The TEVA Defendants have represented that they are only using a single tag. That is Responsive or Not (I often herein refer also to relevant/not relevant, as TEVA has at sometimes referenced relevance rather than responsiveness, and since that is the touchstone of this process).

They are not indicating to the system the weight of the relevance nor are they indicating the topic of relevance.

Accordingly, a perfect score of 1.0 **only** will indicate that the machine finds the document to be in high correlation with other documents **currently** marked as relevant.

It does **NOT** bear any indication on whether the document is actually relevant, on which topic it is relevant, nor its degree of relevance with respect to that topic.

A document with a perfect score of 1.0 may only be **marginally** relevant to the litigation, but well matched to many prior reviewed, but marginally relevant documents.

Conversely, a document with a score of 0.0002 might be **extraordinarily relevant** to the litigation, but just different or new in comparison to the rest of the documents that were marked relevant, or it may be very similar to documents previously marked as not relevant.

Let's take three concrete examples.

(1) A testing report showing NDMA contamination

(2) An email addressing NDMA contamination

(3) A chromatogram showing NDMA contamination

A testing report showing NDMA contamination is not merely relevant, but substantially relevant to the issues in the litigation. However, suppose that reviewers have already marked 1000 emails with attached testing reports that show test results across many disparate items as not relevant. The system has effectively been “trained” that emails transmitting testing reports are generally not relevant. So, when the system examines this next testing report, where there’s a singular notation of NDMA in the listed chemicals found, it is not unreasonable to expect that the system, even a very intelligent system, will assign a predictive score close to zero. Thus, because irrelevant testing reports are common, an email referencing a highly relevant report may get a low score, let’s say of 0.0002. I set forth a non-zero score because there is some textual difference from prior marked documents. This means this highly relevant report **AND THOSE REPORTS LIKE THIS ONE** will be deprioritized **even if there are multiple reports similar to this one, but the system was just trained that most reports are irrelevant**.

Let’s move on to my second example, an early email on NDMA contamination. Let’s say that reviewers have reviewed many emails, marking many relevant for many different reasons, but the next email is the very first email on NDMA contamination. The machine’s algorithm will naturally not recognize a clear pattern between this email and the previous ones marked relevant since it has been “trained” on the first emails marked as relevant, which are numerous at this point. So, this email on NDMA contamination will naturally be assigned a lower predictive score closer to zero than one since it may not be a good “match” linguistically to what was previously marked. The score assigned here does not accurately indicate the extent to which this email is relevant. Let’s suppose the machine gives it a score of 0.25. **BUT worse yet the bias persists.** Let’s suppose even though this email has a lower score, it is randomly inserted into the next review queue and reviewed. Let’s even suppose the reviewer marks this email as relevant. Even if the reviewer marks this email as relevant, since this is only 1 email out of let’s suppose 10,000 marked relevant, **the next email on NDMA contamination will still get a score closer to zero than one**. In fact, there is no

guarantee that the next email on NDMA contamination will even get a higher score than the first. Rare items are thus biased in their predictive scores since even the most intelligent system can only base its predictions using previous patterns. But clearly the most relevant items are by definition going to be relatively fewer in number.

TEVA included a footnote that the CMML system is scoring based upon the analysis of the **textual content of a document**. Even if that system worked perfectly, there are multiple categories of documents where the textual content is imperfect or wholly absent. Documents where there is an image / graphic that is highly relevant, hard copy documents that were OCR'd, audio and video files, and many more could have been picked up by the search terms because of their original file name, folder path they were in, or a portion that did get properly OCR'd, yet given very low predictive scores by the CAL software.

Let's take my third example, a chromatogram attached to an email that shows an NDMA spike. Chromatograms are expressed as images. As such they are effectively invisible to TEVA's vendor's selected CAL software.

Again, the scoring has no correlation to how relevant a document is. The best that can be said is that the CMML system will show reviewers sets of documents more likely to be similar to documents they have previously indicated were relevant. It is a misread to impute to the score any meaning other than whether the machine is saying a document matches the existing marked relevant set.

Unfortunately, reviewers who see scores may be biased by them. A reviewer who sees the computer's opinion of relevancy is 0.0002 is only human. This bias becomes self-reinforcing since if reviewers see a consecutive series of low scored documents that they agree are irrelevant, they will begin to "trust" the computer's judgment. The weight of marking rarer documents as irrelevant has a much stronger adverse effect than the positive effect of marking a rare document relevant in a CAL review. As the corpus of relevant documents increase, when a highly relevant document is mistakenly marked as irrelevant, that "trains" the system that that variation is likely irrelevant.

So, counsel's and TEVA's representation that their CMML / CAL process will automatically prioritize the production of **the most relevant** documents does not stand the weight of even a cursory examination.

Clearly, the only prioritization is of responsive documents in the already narrowed set of documents selected by use of the narrowed search terms, and based solely on the definition of responsiveness that TEVA refused to disclose during the meet and confer call, and in my opinion there is no reasonable basis to assert that the prioritized documents are **the most relevant**. In fact, as set forth above, some of the most relevant documents are unlikely to be prioritized (and due to bias may even potentially be excluded) than if the ordered linear review, with targeted prioritization, is implemented. Moreover, since search terms were already used to substantially narrow the document set, some quantity of documents that were excluded would have been captured by the CAL process if it were applied to the entire document set.

ANY SYSTEM IS ONLY AS GOOD AS ITS INPUTS

In some ways, these complex CAL systems can be thought of in very simple terms, as akin to a calculator.

In elementary school, when one is first introduced to a calculator, it appears to be a magical tool. For any simple numerical problem, you just enter in the values and operations, and *voila* out appears an answer.

For the slightly more complex operations, such as square roots, it's even more magical that it can provide an answer so quickly.

Unfortunately, one of the first hard lessons you learn with a calculator is that merely getting an answer does not mean you got the **right** answer.

At the July 15, 2020 hearing, defense counsel talked at length and at multiple points about the intelligence of TEVA's machine.

TEVA's machine is like any other. It is only as good as its inputs. It may be a very sophisticated calculator, but Plaintiffs cannot judge or, more important, ensure its precision and accuracy without examining the inputs along with the outputs. Moreover, as stated herein, the CAL system should be applied to the entire collected document set for production purposes, not *ex post facto* on the narrowed set created by prior application of search terms.

CAL SYSTEMS AT THEIR HEART ARE PATTERN RECOGNIZERS AND CODIFIERS OF BIAS

These CAL systems are much more sophisticated than a calculator, but conceptually and practically, they can be thought of as sophisticated pattern recognizers.

There is another way to think of CAL systems. They codify bias.

I do not mean to impute any negative connotation to the word bias. What I mean to say is that the bias of early reviewers naturally becomes self-reinforcing in any CAL system.

The conceit of CAL is that it attempts to standardize those individual biases across the whole set of reviewers.

With CAL, unlike in a linear review, decisions made by early reviewers strongly influence which documents are shown next. Even if scores are not shown to reviewers, the very fact that a reviewer gets series of mostly relevant or mostly irrelevant documents in effect "trains" the reviewer. An irrelevant document in a sea of relevant documents is more likely to be marked as relevant. Conversely, and of greater concern as to bias against the Plaintiffs,

relevant documents, even highly relevant, are less likely to be detected as relevant by a reviewer if they are presented in a sea of documents deemed irrelevant.

Ultimately, these observations simply demonstrate what is apparent. Without the full engagement of the Plaintiffs from the outset, which would have been a time consuming process, and would require ongoing engagement on a going forward basis in order to combat and quality check the biases, the production resulting from application of CAL to the truncated set of documents utilized provides no reasonable assurance of fairness or accuracy.

A HAY STACK WITH MANY NEEDLES

Review is not uncommonly compared to as finding the needles in the haystack. The hay and needles represent the documents collected. The needles are the documents relevant to the litigation.

But as discussed before, not every needle is the same. Let's say that the sharpest needles represent the documents **most relevant** to the litigation whereas blunt needles are relevant, but not necessarily the most relevant.

Our search term process, through the multiple rounds, was designed to weed out hay. Of course, in reducing the size of the pile, we accepted that we probably left out some needles, but we agreed that there was likely a greater concentration of needles in what we were retaining for review.

Since we do not know otherwise and as we select batches to examine at random, we have to assume that the needles are evenly distributed in the hay.

In a linear review, a random bunch of hay is grabbed and examined. The likelihood of that hay containing needles both blunt and sharp is equally likely for any other batch. So, while there is hay to sort through, there is no prejudice. We all accept that in reviews of large

volumes of documents with limited time, some needles might be missed. But the important part is that there is no inherent bias in the process. And, the search terms were applied to the entire stack of hay.

CAL is very different, especially as proposed by TEVA to be only be applied to the hay and needles that were identified by the search terms – excluding a vast expanse of hay and by extension, needles, that should have been searched from the outset.

In a CAL simulated review process, the first few batches of hay are random. Reviewers pull out the needles they see in these batches.

BUT since blunt needles are much more common, statistically, there will be a lot of blunt needles pulled and perhaps a few sharp needles.

The problem compounds in the subsequent batches.

As the system is "trained," it recognizes that the blunt needles are needles that are being identified and pulled out of the hay. So, the next batches are weighted to pull more blunt needles and less other, the other being sharp needles and hay.

After some time, all the blunt needles have been pulled out of the stack. Now, when new batches are examined, the system preferences hay that looks like blunt needles, not necessarily the sharp needles because there is likely more hay that resembles needles in the remaining stack than actual needles. So, reviewers start complaining about the waste of time going through all of this hay that's just hay. They may find some needles, but not think it worth the effort.

The problem is that the sharpest needles are likely, by simple statistical inference, to remain in the stack. Pulling samples from the remaining stack will just validate the review that all of the blunt needles have been produced.

Plaintiffs and Defendants have different perspectives. Plaintiffs believe sharp needles exist and are often rare. Defendants believe they do not exist. The purpose of the review is to identify for production the documents that should be produced.

Plaintiffs and Defendants have already worked to refine the search terms, in an effort to identify the blunt and sharp needles, and the narrowed the search terms extensively, and made other reasonable concessions to efficiency. For example, when words were showing up in email footers, we worked collaboratively to exclude those hits. When inspectors' names were common, *e.g.* Patel, we worked to exclude those hits. Should Defendants during their review approach Plaintiffs on a category of documents such as company newsletters that are challenging, Plaintiffs have already told Defendants they would work with Defendants to fashion a mutually agreeable solution. The proposed application of CAL to the already heavily narrowed set of documents was never contemplated, and if it was known that this methodology would be used there would not have been a limitation of the part of the haystack to be searched – the system is built to search the entire hay stack.

If efficiency and saving time and money is truly the goal, using CAL to identify voluminous sets of documents least likely to be relevant, for example, news clipping services that may have gotten unintentionally captured by search terms and could be narrowed through additional searches before reviewing them all, would be more effective. As in the classical solution to burning down the haystack to find the needles as opposed to building ever more complicated needle finding machines, my point is that there are other ways of significantly cutting down the review burden that are less likely to introduce prejudice.

MORE COMPLEX THAN A HAYSTACK

In complex litigation, we do not simply have hay, blunt and sharp needles. Not only are documents more or less relevant, but they are more or less relevant with respect to a large

number of topics. For example, in this litigation, a document could be more or less relevant to what a company told regulators versus what they knew internally.

Instead of a haystack, picture a forest of documents. Surveying the forest, we want to distinguish all of the wildlife from the trees. It is too simplistic to just categorize into tree or not. We want to categorize into birds, insects, larger creatures, and unexpected visitors.

Unfortunately, if we configured a CAL system in the same manner as TEVA's proposal, it would tend to get "stuck" on a particular topic until it ran out of that topic. If sparrows are common in our forest, reviewers would be shown pictures of every sparrow until they were exhausted. Then the CAL system would get "stuck" on the next random relevant category that grabbed its attention. It would become fixated and refuse to show pictures of anything else.

Why? Because in TEVA's configuration, the system has no inputs on topic, no assurance of contextual diversity, and no feedback on degree of relevance. It can only codify bias based upon what it is told.

The natural consequence is an inverted production where the least interesting relevant documents are produced first and in great quantity. In our forest survey, we want to learn about the rarer migratory animals in each specific category, not just get a dump of pictures of ants, sparrows, and squirrels.

In our more complex haystack, an early review cutoff deprives the receiving party of the most important information across multiple different topics.

In this litigation, there are numerous key topics fundamentally distinct from one another, represented by different and overlapping sets of documents. A number of these topics have yet to be identified and will only become apparent during document reviews and depositions.

EXAMPLES OF UNEXPECTED BIAS IN REAL WORLD AI SYSTEMS

Although machine learning systems (AI) are expected to be neutral, fair, and intelligent, there are plenty of ways that bias can be unwittingly introduced, negating the perceived impartiality of the output.

As Karen Hoa wrote in her review of AI Bias in MIT Technology Review last February, https://www.technologyreview.com/2019/02/04/137602/this-is-how-ai-bias-really-happensand-why-its-so-hard-to-fix/, bias can be introduced during framing the problem, collecting the data, or preparing the data. It can be difficult to detect:

> "The introduction of bias isn't always obvious during a model's construction because you may not realize the downstream impacts of your data and choices until much later. Once you do, it's hard to retroactively identify where that bias came from and then figure out how to get rid of it. In Amazon's case, when the engineers initially discovered that its tool was penalizing female candidates, they reprogrammed it to ignore explicitly gendered words like "women's." They soon discovered that the revised system was still picking up on *implicitly* gendered words—verbs that were highly correlated with men over women, such as "executed" and "captured"—and using that to make its decisions."

When Amazon used AI in their hiring, they discovered afterwards it was biased against women. When NVIDA developed an AI to unblur faces, they discovered the system was outputting faces that transformed the faces of people of color into faces that looked distinctly white, https://www.theverge.com/21298762/face-depixelizer-ai-machine-learning-tool-pulse-stylegan-obama-bias. Similarly unexpected bias has been documented even with large sets of training data, such as the 14MM images tagged as part of the ImageNet data set where matching on neutral terms such as "programmer" yielded images of white men: https://www.wired.com/story/ai-biased-how-scientists-trying-fix/. Facial recognition algorithms routinely struggle with bias: https://www.wired.com/story/best-algorithms-struggle-recognize-black-faces-equally/. Even Google has struggled to eliminate contextual bias in simple search prompts: https://www.wired.com/story/google-autocomplete-vile-suggestions/. This has also impacted the Courts as the well documented bias in the COMPAS

system used to predict rates of recidivism: https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm.

Accordingly and unsurprisingly, in August 2019, the American Bar Association, recognizing the challenges of bias in AI system passed Resolution 112: (https://www.americanbar.org/content/dam/aba/images/news/2019/08/am-hod-resolutions/112.pdf)

> RESOLVED, That the American Bar Association urges courts and lawyers to address the emerging ethical and legal issues related to the usage of artificial intelligence ("AI") in the practice of law including: (1) bias, explainability, and transparency of automated decisions made by AI; (2) ethical and beneficial usage of AI; and (3) controls and oversight of AI and the vendors that provide AI.

The resolution which was adopted at the Annual Meeting and accompanying announcement, https://www.americanbar.org/groups/business_law/publications/committee_newsletters/legal_analytics/2019/201908/ai_law/, together make clear that any adoption or use of TAR/CAL should be done transparently with an eye on bias.

CAL PROPERLY DONE

As with any tool, there are proper and improper applications of CAL systems. As a threshold, as set forth above, the system is designed to be applied to the entire document set. In an effort to address a fair application of this system, and to respond to the Court's inquiry regarding necessary tagging, it is my opinion that in a proper application of CAL in this or other similar complex litigations, the producing party must:

(1) use multiple tags / vectors for relevance as negotiated with the receiving party instead of a universal single relevant/not relevant tag, i.e. relevant because... In the absence of a transparent meet and confirm with TEVA's counsel and vendor, it is difficult at this point to identify with precision all of the additional tags that should be used. Having said that, at a

minimum, in this litigation, high level concepts such as contamination, chemistry issues, quality control, quality assurance, sales/financial, marketing, cGMP, adverse events, inspections, regulatory, chromatography and mass spectrometry testing are just some of the topics that would need to be coded in order for a CAL system to be able to properly prioritize and mark for relevance.

(2) use at least one tag to indicate the relative weight of relevance, e.g. clear, probable, possible, neutral, contextual

(3) implement prioritization through targeted initial seed sets

a) run searches for the most relevant documents; this has to be done categorically for each category
b) mark the relevant documents in each category (using senior reviewers)
c) feed back into the system the relevant documents (excluding ones marked irrelevant) to train it in each categorization
d) use senior reviewers to conduct quality assurance
e) share initial seed sets with the receiving party
f) accept initial seed sets from the receiving party

(4) use random sampling and CAL at continuous points throughout the process to re-examine documents reviewers marked non-relevant

a) reintroduce any documents marked as not relevant in error back into the CAL system
b) rerun CAL to find similar documents marked as not relevant that should have been marked as relevant

(5) do training side by side with the receiving party until the receiving party is in agreement on what constitutes relevant documents

(6) on a separate CAL project, use tagging to weed out document types (ex. vacation requests, janitorial supply invoices, etc.) the producing party believes are categorically not relevant (but these would need to be logged so that if the receiving party disagreed that these documents are not relevant there is transparency and that decision could be challenged)

(7) ensure the introduction of contextual diversity in subsequent review batches; this can be set through appropriate configuration flags

(8) suppress any CAL relevance scoring when reviewers are reviewing a document to avoid introducing bias

(9) set a threshold score above which documents will be automatically produced without any further review where CAL has scored those documents with a mutually accepted high degree of probable relevancy in one or more categories

(10) after reviewing an agreed upon percentage of documents not less than 10%, automatically produce (subject to privilege review) all documents the CAL system has marked as relevant above a certain threshold point

(11) not automatically apply CAL training to subsequent custodians or sets of documents that are expected to be significantly different from the prior reviewed corpus

(12) determine *a priori* whether the set of collected documents is large enough and diverse enough to support CAL or otherwise inappropriate for CAL; not every set of documents would be appropriately reviewed by CAL

(13) abstain from using CAL for certain categories of documents where relevance is near impossible to recognize outside of context or where CAL is ill suited to do a proper linguistic analysis, ex: very short emails, spreadsheets, chromatographs, other images, audio files, video files, OCR'd files

(14) conduct extensive targeted sampling (with full disclosure of all sample documents to the receiving party) before proposing a discard point if the producing party is not using CAL to only prioritize the review (in the situation here, where the initial groundwork has not been done, all documents targeted for exclusion should be manually reviewed).

UNWORKABLE AT THIS POINT IN THE PROCESS

During the July 15th, 2000 hearing with the Court, Teva's counsel cited to the TAR/CAL protocol that was adopted In Re Broiler Chicken Antitrust Litigation, Case No. 1:16-cv-08637 N.D. Illinois.

While a discussion of CAL/CMML at the outset would have given the Parties enough time to properly implement joint reviews and an overall protocol for review, at this point in the process with the production of documents underway, establishing a workable CAL process that is fair and equitable and that will result in timely, fully prioritized productions, is not reasonably achievable, and the effort would impose a significant distraction on the Plaintiffs. The sudden suggestion by ZHP and Mylan that they might elect to use CAL only would multiply the difficulty.

The protocol in the Broiler Chicken Antitrust litigation was set out right at the beginning of the process, not foisted on at the last minute. The TAR/CAL system to be used had to be disclosed up front, prior to any search terms being negotiated. The Plaintiffs could opt out custodians from the search terms and the TAR/CAL process and subject them solely to the TAR/CAL process. Those threshold steps did not occur here, and cannot occur here, due to the belated revelation of Teva's intent to do this. Unfortunately, the validation process laid out in that protocol was very limited and inadequate for a litigation of this size and complexity. A great deal of work would need to be focused on establishing an appropriate validation process for this litigation.

The TEVA Defendants are pioneering CAL review systems within this MDL. They intend to do so unilaterally. They likely intend to use these systems to support future arguments for thresholds of scores below which they will not review documents or will offer cursory examination.

Contrary to any assertion that the system is currently "learning" intelligently, the best evidence is that the system is preferencing more numerous documents over documents that are rarer or different, but no less relevant or critical to proving this case.

This is not a matter of sitting down once with TEVA to work out a protocol. Due to the nature of CAL, Plaintiffs would have to be involved in the actual responsiveness review. Plaintiffs would have to be intimately involved in the training of the system. That is not possible at this point.

It is clear that Defendants are beta testing this system. This is a complicated system with complex side effects. Errors and missteps are costly both in time and money to fix.

At this point, doing CAL properly would add a potentially drastic amount of time to the review and require collaborative retraining of the existing CAL system. Again, this would be a significant distraction where the Plaintiffs are now ingesting document productions on a rolling basis from numerous defendants.

INFEASIBILITY OF VALIDATING POST PRODUCTION

It is infeasible for TEVA to ask the Plaintiffs meaningfully to evaluate a CAL process post the initial productions.

First, the initial productions will, by definition, contain relevant documents. How are Plaintiffs to determine which categories of relevant or prioritized documents are missing,

especially if the Defendants simply argue that they do not exist? There is no way for Plaintiffs to know.

PLAINTIFFS ARE UNDULY PREJUDICED EVEN IF EVERY DOCUMENT IS ULTIMATELY REVIEWED

For the reasons outlined before, rarer but highly relevant documents, “diamonds,” will likely be pushed to the last of the rolling productions. That effectively negates the intent of a rolling production, and absolutely negates the prioritization by Plaintiffs. This will result in Plaintiffs having to search a potentially large production to identify documents that in many cases should have been produced at the very outset.

AFFIRMATION

Pursuant to 28 U.S.C. § 1746, I declare under penalty of perjury of the laws of the United States of America that the foregoing is true and correct.

Executed on this 24th day of July, 2020, at Forest Hills, NY.

Digitally signed by Jonathan Jaffe
DN: cn=Jonathan Jaffe, o=www.its-your-internet.com, ou, email=jjaffe@its-your-internet.com, c=US
Date: 2020.07.24 13:59:29 -04'00'

Jonathan Jaffe